

[illegible]

5

10

2. Description of Related Art

25

30

35 Figure 2 illustrates a server cluster. Figure 2 shows client computers 10 connected to a switch 20, which together form a local area network (LAN). The switch connects the local area network to the server cluster, which in this example comprises five server computers 30. Figure 2 further illustrates mass storage units 60 connected to

In general, there are two basic methods to construct server and gateway clusters in a TCP/IP network. According to the first basic method, each server or gateway has a distinct IP address, and an external device redirects the requests or packets to different servers or gateways. The external device has a certain set of rules for controlling the redirection of data packets. For example, the external device can redirect data packets on the basis of current load on the cluster devices. In the example of figures 1 and 2, the external device can be the switch 20 connecting the LAN to the cluster. The main drawback of this solution is asymmetric routing: because routing in IP networks is usually based only on the destination IP address, outgoing and incoming packets can be routed to different gateways. This is a problem when the gateways also serve as firewalls, since firewalls typically do not accept any incoming messages, which are not a reply to a previous outgoing request. When the incoming packets are routed to a second gateway which is a different gateway than the first gateway which sent the original outgoing request, they will be discarded since the second gateway does not have knowledge about the previous request.

According to the second basic method, all units of the cluster share a common IP address and MAC (media access control) address, and therefore receive all traffic in both directions. The units of the cluster have filter functions at their inputs and outputs, which according to certain rules allow the processing of a packet by only one of the gateways, while the others ignore the packet. The rules are constructed in such a way that the same gateway processes the packets in both directions of a single connection, which allows normal firewall operation. The benefit of this solution is transparency and scalability. New servers or gateways can be added to the cluster without changing routing, i.e. unnoticed by the clients.

However, even the second basic method has its drawbacks. Connecting multiple servers or gateways sharing a common MAC address to a switch requires usually usage of a multicast MAC address for the cluster, i.e. a hardware layer address which is recognized and processed by the network interface of every unit of the cluster. The network interfaces of the gateways always have individual MAC addresses as well, but in this solution, the interfaces are arranged to recognize the multicast MAC address selected for the cluster. The IP address of the server or gateway cluster is mapped to a multicast MAC address in the ARP (Address

Resolution Protocol) tables of the nodes of the IP network. When a client wishes to send a data packet to the cluster, it addresses the packet to the IP address of the cluster. The network translates the IP address to a MAC address for transportation on the physical layer, and in this case the MAC address is a multicast address. The
5 problem with multicast MAC addresses is that switches typically flood packets having a multicast address as destination to all ports of the switch by default, wherefore the packet ends up in all other subnetworks connected to other ports of the switch. This effect causes an unnecessary loading of other subnetworks connected to other ports of the switch than the gateways/servers of the cluster. This
10 is a severe problem, if the LAN is divided into virtual LAN networks (VLAN).

Virtual LANs (VLAN) can be created using switches to segment networks. A single segment is a broadcast domain, i.e. broadcast, multicast and unknown unicast frames are sent to all stations in the segment. Port-based virtual LAN networks are
15 groups of switch ports and attached segments (subnetworks) which belong to the VLAN.

SUMMARY OF THE INVENTION

20 An object of the invention is to realize a method for transmission of data in an IP packet network, which avoids the problems of the prior art.

The objects are reached by using the IGMP protocol to inform the switch, which of its ports are connected to units of a cluster system, and forwarding those packets
25 with the unicast IP address of the cluster system and the multicast MAC address of the cluster system only to those ports.

The method according to the invention is characterized by that, which is specified in the independent method claim. The cluster system according to the invention is
30 characterized by that, which is specified in the independent claim directed to a cluster system. The dependent claims describe further advantageous embodiments of the invention.

The invention is applicable to gateway clusters as well as server clusters. In the
35 following, the single unit of cluster i.e. a gateway or a server is denoted a cluster unit.

09686273-101000

According to the invention, the cluster units are configured to be members of an IP multicast group specific to the cluster. The switch or switches directly connected to the cluster units are arranged to monitor multicast group membership reports from the cluster units, and therefore obtain knowledge about which ports of the switch or switches are connected to cluster units. Advantageously, the switch or switches may also send membership queries to find out, which ports are connected to members of the cluster multicast group. Consequently, when the switch receives a packet with a multicast MAC address and the unicast IP address of the cluster, the switch sends the packet to only those ports to which cluster units are connected, and not to all ports of the switch as according to the prior art. In the local area network, the unicast IP address of the cluster is mapped to a multicast MAC address in the routing tables in the network. Therefore, when any client sends a packet to the cluster, the switch will receive the packet with a multicast MAC destination address.

For obtaining this functionality, the switch is arranged to recognize the unicast IP address of the cluster and the multicast group address configured for the cluster, and to store into its memory information about which ports are connected to such network nodes, which send group membership reports for the multicast group address configured for the cluster.

According to the IGMP protocol (Internet Group Management Protocol), network nodes send a group membership report when they join a multicast group to announce their presence to multicast routers. The report is sent with an IP destination address equal to the host group address being reported, so that the multicast router can recognize which multicast groups have members in its domain. A multicast router sends an IGMP query at regular intervals to see if any nodes belong to any multicast groups. Network nodes send a group membership report, when they receive a group membership query from the router. According to the invention, this mechanism is used to inform the switch about which ports are connected to cluster units. The IGMP protocol version 1 is defined in the IETF standard document RFC 1112, and IGMP version 2 is defined in the IETF proposed standard document RFC 2236.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention is described in more detail in the following with reference to the accompanying drawings, of which

Figure 2 illustrates a server cluster arrangement according to prior art,

Figure 4 illustrates another method according to an advantageous embodiment of the invention, and

Figure 5 illustrates a cluster system according to an advantageous embodiment of the invention.

Same reference numerals are used for similar entities in the figures.

A. A FIRST GROUP OF ADVANTAGEOUS EMBODIMENTS

Figure 3 illustrates a method according to an advantageous embodiment of the invention. Figure 3 illustrates a method for transmission of data in an IP packet network, which network comprises a cluster of cluster units and a switching unit having a plurality of ports, and in which network the cluster units are connected to a part of the plurality of ports, and which cluster units share a unicast IP address.

According to the present embodiment, the method comprises at least steps, in which

- the cluster units are configured 100 to be members of an IP multicast group specific to the cluster,

- the IGMP protocol is used to obtain information about to which ports of the plurality of ports the cluster units are connected,

- the MAC address of a received IP packet is checked in step 140, and if said MAC address is found to be a multicast MAC address, the IP destination address of said packet is compared 150 to the unicast IP address shared by the cluster units, \

- if the IP destination address of said packet is the same as the unicast IP address, the packet is forwarded 160 to those ports, to which the cluster units were found to be connected, and if it is not, the packet is sent to all ports of the switch.

Figure 4 illustrates as an example how the step 110 of obtaining information is performed according to an advantageous embodiment of the invention. According to this embodiment, the method further comprises at least steps, in which the switching unit receives 114 an IGMP group membership report, the switching unit checks 116, if said report is addressed to said IP multicast group specific to the cluster, and if it is, the switching unit stores 118 into its memory the identifier of the port, through which said report arrived.

Advantageously, the method may as well comprise a step, in which an IGMP group membership query is sent 112 by the switching unit.

In the inventive method, the IGMP protocol version 1 according to RFC 1112 can be used, as well as the IGMP protocol version 2 according to RFC 2236.

B. A SECOND GROUP OF ADVANTAGEOUS EMBODIMENTS

Figure 5 illustrates a cluster system according to an advantageous embodiment of the invention. Figure 5 shows client computers 10 connected to a switch 200. The client computers 10, the switch 200, and form a local area network (LAN). The LAN is connected to an external network 50 through a gateway cluster consisting in this example of five gateways 30. The gateways 30 may for example function as firewall devices. The gateways 30 are connected to the external network 50 through a switch 200 and a router 40. The client computers are in figure 5 divided into three subnetworks A, B, and C. According to the embodiment shown in figure 5, the cluster system has a plurality of cluster units 30 and a switching unit 200, and said cluster units are associated with the same IP unicast address. According to the embodiment, the cluster units are configured to be members of an IP multicast group specific to the cluster system, and the cluster system comprises

- means 240 for observing using the IGMP protocol which ports of the switching unit are connected to the cluster units,
- means 210 in the switching unit for observing the MAC destination address of a packet arriving to the switching unit and for checking if said MAC destination address is a MAC multicast address,
- means 220 in the switching unit for observing the IP destination address of said packet and for comparing said IP destination address to said IP unicast address associated with the cluster units,
- means 230 in the switching unit for forwarding of the packet to those ports whose identifiers were previously stored to said memory means as a response to the finding

that said IP destination address and said IP unicast address are the same and said MAC destination address is a MAC multicast address.

In a further advantageous embodiment of the invention, said means 240 for observing using the IGMP protocol comprise at least

- means 242 in the switching unit for observing IGMP multicast group reports and for checking, if a received IGMP multicast group report is addressed to said IP multicast group specific to the cluster system, and

- means 244 in the switching unit for storing into a memory means an identifier of that port via which said received IGMP multicast group report arrived as a response to finding that said report was addressed to said IP multicast group.

In a still further advantageous embodiment of the invention, the system further comprises means 260 in the switching unit for sending IGMP group membership queries.

In an advantageous embodiment of the invention, said means 210, 220, 230, 240, 242, 244, 250, and 260 are realized using software programs stored in a memory element 270 of the switching unit 200 and executed by a processor 280 of the switching unit 200.

In another advantageous embodiment of the invention, said means 210, 220, 230, 240, 242, 244, 250, and 260 are realized using dedicated logic circuits implemented for example using an ASIC circuit (application specific integrated circuit). Such an implementation of the invention allows the functioning of the switching unit at extremely high speeds.

In an advantageous embodiment of the invention, the cluster units 30 are gateway units 30.

In a further advantageous embodiment of the invention, the cluster units are server units.

C. FURTHER CONSIDERATIONS

The invention has several advantages. For example, the invention enables connecting gateway clusters to switches, which allows the construction of high-speed networks using port-based virtual local area networks.

5

10